

Practical data synthesis for large samples

Gillian M. Raab*, Beata Nowok† and Chris Dibben‡

School of Geosciences, University of Edinburgh, Edinburgh, UK

March 22, 2016

Abstract

We describe results on the creation and use of synthetic data that were derived in the context of a project to make synthetic extracts available for users of the UK Longitudinal Studies. A critical review of existing methods of inference from large synthetic data sets is presented. We introduce new variance estimates for use with large samples of completely synthesised data that do not require them to be generated from the posterior predictive distribution derived from the observed data. We make recommendations on how to synthesise data based on these findings. An example of synthesising data from the Scottish Longitudinal Study is included to illustrate our results.

Keywords— synthetic data, disclosure control, CART, UK Longitudinal Studies

1 Introduction and background

1.1 Synthetic data for disclosure control

National statistics agencies and other groups collect large amounts of information about individuals and organisations. Such data can be used to understand population processes so as to inform policy and planning. The cost of such data can be considerable, both for the collectors and the subjects who provide their data. Because of confidentiality constraints and guarantees to data subjects full access to such data is often restricted to the staff of the collection agencies. Traditionally, data collectors have used anonymization along with simple perturbation methods such as aggregation, top-coding, record-swapping, suppression of sensitive cells or adding random noise, to prevent the identification of data subjects. Advances in computer technology and search techniques have illustrated how such measures may not prevent disclosure [36]. Additionally, examples show that these ad hoc disclosure procedures may compromise the conclusions that can be drawn from such data [19, 60].

Synthetic data, which retains the essential features of the actual data with some or all of the values replaced by simulations from probability distributions, can be a way of overcoming these difficulties. There is now an extensive literature on this topic [6, 10, 13, 25, 39, 40, 43, 44], much of it summarised in the monograph by Drechsler [9], as well as several initiatives, e.g. [56, 57], which make synthesised data available to researchers.

*gillian.raab@ed.ac.uk

†beata.nowok@ed.ac.uk

‡chris.dibben@ed.ac.uk

1.2 Application to the UK Longitudinal Studies

The England and Wales Longitudinal Study (ONS LS) [21] the Scottish Longitudinal Study (SLS) [4] and the Northern Ireland Longitudinal Study (NILS) [37] are rich micro-data sets linking samples from the national census in each country to administrative data (births, deaths, marriages, cancer registrations and other sources) for individuals and families across several decades. Researcher access to the UK Longitudinal Studies (LSs) is highly restricted due to confidentiality and legal constraints. Thus the three LSs have a small number of users compared to other census data products. Synthetic data with no real individuals, but which mimic the real data and preserve the relationships between variables and transitions of individuals over time, could be made available to accredited researchers to analyse on their own computers.

Although the LSs are pseudo-random samples from the censuses, with linked administrative data, people use the data to draw general conclusions about social processes, rather than to infer anything about the census totals. Thus inferences from the LSs generally assume that the data are samples from a hyper-population model that might have generated the data. Methods appropriate for simple random sampling (SRS) are generally used, except for analyses that involve clustering, e.g. in families or by geographic area, when hierarchical models may be used.

Other projects that make synthetic data available have been based around a single data set, so that synthesised data sets are created once and extracts from them are supplied to users. Every user of the LSs has a customised, linked data set made available to them. Thus a new synthesis is required for every researcher. To make it possible for LS support staff to provide synthesised extracts for researchers we have developed the **R** package *synthpop* [35] which is now available free to any user.¹ As we implemented different proposed methods of inference from synthetic data, we were led to re-evaluate the conditions under which each is valid. We have also derived new variance estimates which can be used under certain conditions and which can be calculated from a single synthetic data set.

1.3 Review and critique of methodology for inference from synthetic data

Rubin's original proposal for synthetic data [51] saw it as an example of multiple imputation (MI). The first paper to provide detailed methodology for making inferences from synthetic data [39] followed this approach. Synthesis was assumed to proceed by imputing all the unobserved members of a population and then selecting a synthetic sample from each synthetic population. Because imputation is involved, it was considered mandatory to generate the synthetic populations from their posterior predictive distribution (PPDs), given the observed data. The variance of estimates from synthetic data differ from those for MI by replacing the addition of the term for within-imputation variance by its subtraction. The authors ([39] section 2.2) attribute this difference to the sampling for each synthetic sample from its population and this explanation is expanded in [47], section 3.1. This explanation is correct in that the between-synthesis variance includes a contribution from sampling from the synthesis distribution, but it is not the underlying reason for the negative sign in the variance estimate. We will show here that results for inference from synthetic data can be derived without considering it as imputation at all. The negative sign arises from the need to get the contribution to the estimate from the synthetic data from the uncertainty about the population parameters if inference were from the observed data. When the estimate would be calculated in a different manner from the observed and synthetic data, we need multiple syntheses sampled from the PPD of Y_{obs} so that this uncertainty contributes to the between synthesis variation. The variance contribution from the observed data can then be obtained by subtracting the within-synthesis variance from the between-synthesis variance. For the case when the same estimation methods would be used for the observed and synthetic data, the contribution from the fit to the observed data can be estimated from the within-synthesis variation. This leads to new variance estimates that can be calculated from a single synthetic data set without sampling from the PPD.

If data consist of $(Y|X)$ and inference is required for a statistic Q from the synthesised data,

¹Its initial development was part of the SYLLS (Synthetic Data Estimation for UK Longitudinal Studies) project <http://www.lscs.ac.uk/projects/synthetic-data-estimation-for-uk-longitudinal-studies/> funded by the UK Economic and Social Research Council.

conditional on unchanged variables X , then the variance estimator derived in [39] and our new variance estimators require that all of the observed data Y are replaced by synthetic values. Reiter [41] derived results for making inferences when only part of Y is synthesised and Reiter and Kinney [45] pointed out that sampling from the PPD was not needed for what they term “partially synthetic” data. Their examples include cases where all the data are synthesised. They distinguish fully synthetic data, which requires sampling from the PPD, from partially synthetic data, which does not, by the fact that no new units are introduced in partially synthetic data so that synthetic populations sampled from the PPD are not needed.

In this paper we will show that data synthesis, as usually carried out, does not require the creation of synthetic populations, except in the sense of defining them in terms of some appropriate model. This is true whether or not the synthetic data includes new units that were not part of the observed data. Thus Reiter and Kinney’s distinction between fully and partially synthetic data is not meaningful. The real reason why PPD is not required for “partially synthetic” data is that the same estimators are used to analyse the observed and synthetic data. This will always be true when some observed data are retained and it is also the case for the examples of “partial synthesis” cited in the literature. To avoid confusion with the previous terminology we will use the term “completely synthesised” when all of Y is replaced by synthetic data and “incompletely synthesised” when some of the original Y are retained. Although PPD is not required for incompletely synthesised data, multiple syntheses are still needed because the between-synthesis variance will no longer include all of the contribution from the fit to the observed data.

In most implementations of synthetic data generation, including *synthpop*, the joint distribution is defined and synthesised in terms of a series of conditional distributions. This approach has been used extensively in implementations of MI, e.g. [38, 49, 59]. In MI every variable with missing data is predicted from its conditional distribution and the procedure is repeated sequentially until convergence. This has led to concerns about its validity when the conditionals that define each variable are not compatible [58]. This concern does not apply to synthetic data because the joint distribution is built up from a single set of conditionals and there are no iterations that need to converge to a joint distribution. In our experience synthesising data from conditional distributions works well and allows great flexibility to produce plausible synthetic data that mimics the structure of real survey or administrative data. A more important requirement underlies all synthetic data estimation. This is the assumption that the real data were generated from the same model that was used for synthesis. We will refer to this as the Assumed Synthesising Distribution assumption, or the ASD for short. The results derived in this paper only apply for large samples, making the same asymptotic assumptions as in the other literature cited, and where the ASD is met.

1.4 Structure of this paper

The structure of this paper is as follows. Section 2 presents our results on inference from synthetic data. This is first developed for the case when methods appropriate for simple random sampling can be used to obtain estimates for both the observed and synthetic data. Completely synthesised data are considered first, our new variance estimators are introduced and their properties confirmed by a simulation. Results for incompletely synthesised data are presented in section 2.3 and the extension to methods for complex samples in 2.4. We conclude this section by summarising our recommendations for practice based on these results. Section 3 reports on our practical experiences of creating synthetic data using *synthpop* and section 4 briefly reviews aspects of the disclosure risk and utility of synthetic data. An example of synthesising data from the SLS is presented in section 5 and compares the different variance estimators discussed in section 2. The final section presents a summary and proposed future directions for our work.

2 Inference from synthetic data

2.1 Completely synthesised data

Here we provide a derivation of variance estimator (T_M) [39], without the requirement to create synthetic populations. The analyst wishes to make inferences about a quantity Q , given only the synthetic data. We can write the required posterior distribution as

$$Pr(Q|X, D_{syn}) = \int Pr(Q|X, Y_{obs}, D_{syn}) Pr(D_{syn}|X, Y_{obs}) dY_{obs} \quad (1)$$

where Y_{obs} is the observed data, D_{syn} the synthetic data and inference is conditional on the unchanged variables X . This compares to the equivalent expression for MI ([50], equation 3.2.2, p.82) where integration over the unobserved elements of Y requires sampling from their PPD. In equation (1) D_{syn} are not unobserved data, but are random functions of Y_{obs} . For clarity of exposition we drop the conditioning on X so that (1) becomes

$$Pr(Q|D_{syn}) = \int Pr(Q|Y_{obs}) Pr(D_{syn}|Y_{obs}) dY_{obs} \quad (2)$$

We assume that D_{syn} consists of M synthetic data sets ($Y_1, Y_2, \dots, Y_l, \dots, Y_M$) each generated by SRS from the PPD of Y given Y_{obs} , assumed to be generated from some distribution $f(Y|\theta)$. From each Y_l we calculate $q^{(l)}$ and $v^{(l)}$, the estimate and variance of Q calculated as if the Y_l were the observed data. We aim to make inferences for Q based on

$$\bar{q}_M = \sum_{l=1}^M q^{(l)}/M, \quad \bar{v}_M = \sum_{l=1}^M v^{(l)}/M \quad \text{and} \quad b_M = \sum_{l=1}^M (q^{(l)} - \bar{q}_M)^2 / (M - 1)$$

Paralleling [39], we make the following asymptotic assumptions.

1. Sample sizes are large enough to permit Normal approximations to the posteriors for all estimates of Q with locally flat priors.
2. The estimate \hat{Q} of Q from Y_{obs} is unbiased with variance U , estimated by \hat{U} and the sampling variability of U is negligible relative to that of \hat{Q} , so we can write $\hat{U} \approx U$.
3. If $\hat{\theta}^{(l)}$ is the fitted value of θ used to generate Y_l , then the expectation of $q^{(l)}$ conditional on $\hat{\theta}^{(l)}$ is $Q^{(l)}$ with variance estimated by $V^{(l)}$, again conditional on $\hat{\theta}^{(l)}$, and the variance of $V^{(l)}$ across the M syntheses is negligible so we can write $V^{(l)} \approx V$ for all l .
4. The sampling variance estimate $v^{(l)}$ is unbiased for $V^{(l)}$ and the sampling variance of $v^{(l)}$ is negligible. Thus by 3 we have $v^{(l)} \approx V$ and $\bar{v}_M \approx V$.

Note that in assumption 3 our conditional expectation $Q^{(l)}$ and its estimated variance $V^{(l)}$ replaces the use of these symbols as the properties of a synthetic population as defined in [39]. We also require the assumption that the population size is large enough for the finite population correction factor (fpc) to be ignored. This condition was not mentioned explicitly in [39], but their methods require it. Note that this is true even when whole synthetic populations are generated. The fpc will not reduce the variance of samples from synthetic populations because the totals of elements contributing to estimators will not be constant across populations and will not usually correspond to the totals for the population from which the observed data were sampled. We can write equation (2) in terms of conditional distributions as

$$Pr(Q|D_{syn}) = \int \left\{ \int \left[\int Pr(\bar{q}_M|\bar{Q}_M, \hat{Q}, Q, V, U) d\bar{q}_M \right] Pr(\bar{Q}_M|\hat{Q}, Q, V, U) d\bar{Q}_M \right\} Pr(Q|\hat{Q}, V, U) d\hat{Q}, \quad (3)$$

where $\bar{Q}_M = \sum_{l=1}^M Q^{(l)}/M$. Conditions 2 and 4 justify the omission of integration with respect to U and V in (3) and the replacement of each by their estimates for the final inferences. We can replace

Y_{obs} in (2) by \hat{Q} because conditions 1 and 2 mean that \hat{Q} is a sufficient statistic for Q given Y_{obs} . Now conditions 1 to 4 allow the following approximations, omitting the conditioning on Q ,

$$\bar{q}_M|\bar{Q}_M, \hat{Q}, V \sim N(\bar{Q}_M, V/M), \bar{Q}_M|\hat{Q}, U \sim N(Q, U/M) \text{ and } Q|\hat{Q}, U \sim N(Q, U),$$

where the first two results follow from $q^{(l)}|Q^{(l)}, \hat{Q}, V \sim N(Q^{(l)}, V)$ and $Q^{(l)}|\hat{Q}, U \sim N(Q, U)$ and the conditional independence of $(q^{(l)} - Q^{(l)})$ and $(Q^{(l)} - \hat{Q})$ for different l .

From (3) we have $E(Q|D_{syn}) = E(E(\bar{q}_M|\bar{Q}_M)|\hat{Q}) = Q$ and

$$var(Q|D_{syn}) = V/M + U/M + U = U + (V + U)/M, \quad (4)$$

where the expression for the variance comes from sequential use of the expression for the variance from conditional distributions as the expectation of the conditional variance plus the variance of its expectation. In order to make use of this we need to replace U and V by unbiased estimates. We have \bar{v}_M for V from assumption 4, but for U we need to consider the expectation of b_M ,

$$E(b_M) = E\left\{\sum_{l=1}^M [(q^{(l)} - Q^{(l)}) - (\bar{q}_M - \bar{Q}_M)]^2 / (M-1)\right\} + E\left[\sum_{l=1}^M (Q^{(l)} - \bar{Q}_M)^2 / (M-1)\right] = V + U \quad (5)$$

so that $b_M - \bar{v}_M$ has expectation U and substituting these expectations into (4) we get the asymptotically unbiased variance estimator

$$T_M = b_M(1 + 1/M) - \bar{v}_M. \quad (6)$$

This estimator requires that the synthetic data is generated by SRS, but the real data can come from a complex sample with appropriate methods of inference (e.g. as described in [30]) used in the estimation from Y_{obs} . However, if inference for Y_{obs} has also been carried out assuming SRS, further simplifications of the variance estimator is possible, as described in the next section.

2.2 New simple variance estimators for completely synthesised data

When Q is estimated from Y_{obs} by methods appropriate for simple random sampling and the sample sizes for the real and synthetic data are the same, we can see from assumptions 3 above that $\hat{U} \approx V$, since $\hat{V}^{(l)} \approx V$, and the assumption that $V^{(l)}$ does not vary across synthetic data sets, implies that it will also be the same for the real data. Thus we can replace U by V in (4) and obtain the simple variance estimator

$$T_{s(PPD)} = \bar{v}_M(1 + 2/M). \quad (7)$$

Since (7) does not include U , we no longer need to use equation (5) to get the expectation of U so that multiple syntheses and sampling from the PPD of Y are no longer needed. If we replace sampling from the PPD of Y with generating synthetic data from $f(Y|\hat{\theta})$, then $Q^{(l)}$ will be replaced by \hat{Q} for all l , and the central term in equation (4) drops out to give a simple variance estimator for synthesis without PPD

$$T_s = \bar{v}_M(1 + 1/M). \quad (8)$$

In the more general case when the sample size for the observed data is n and that for the synthetic data is k , equations (7) and (8) become

$$T_{s(PPD)} = \bar{v}_M(k/n + (1 + k/n)/M) \quad (9)$$

and

$$T_s = \bar{v}_M(k/n + 1/M) \quad (10)$$

Note that the estimators $T_{s(PPD)}$ and T_s can be computed from a single synthetic data set with $M = 1$.

The properties of these new estimators were confirmed by simulations one of which is described in A.1. The new estimators were shown to be unbiased and to give the correct coverage estimates. The precision of the new variance estimates is many times better than that of T_M , which can give very unsatisfactory results for small M . This simulation also evaluates the estimator T_p , which is discussed in section 2.3.

Klein and Sinha [26, 27, 28] have derived exact finite-sample results when the observed data are estimated by likelihood methods assuming SRS. Their examples include estimation of the mean for univariate and multivariate Normal distributions and linear regression, with the dependent variable synthesised but the predictors unchanged. These would all be considered as completely synthesised, since for the linear regression inference is conditional on the unchanged X variables. In each case estimates are derived both with and without sampling from the PPD, with the term “plug-in sampling” used for the latter. They note that both estimators can be used for a single multiply imputed data set. For all of their examples their estimates converge for large samples to T_s for plug-in sampling or to $T_{s(PPD)}$ for sampling from the PPD.

2.3 Incompletely synthesised data

When some real data remain in D_{syn} , $(Q^{(l)} - \hat{Q})$ will no longer have variance U , so the estimators T_M and T_s will no longer be unbiased. Reiter [41] demonstrates this and derives a new estimator T_p for incompletely synthesised data, and illustrates the important point that the synthetic data must be based on models fitted only to the data that are to be replaced. The development is similar to our section 2.1. To reproduce it we need to introduce an expression for $E[(Q^{(l)} - \hat{Q})^2] = W$, where we expect $W < U$, since only part of the data are replaced. For sampling from the PPD of the part of the data to be synthesised equations (4) and (5) become

$$var(Q|D_{syn}) = U + (W + V)/M \quad (11)$$

and

$$E(b_M) = V + W \quad (12)$$

and, as before, \bar{v}_M has expectation V . When some observed data remain in the synthetic data, its structure must be the same as that of the observed data, so the analyst with synthetic data will use the same methods as if the observed data were available. Thus we have $U \approx V$, as for the new variance estimators in section 2.2. Substituting the unbiased estimates into (11) we obtain

$$T_p = \bar{v}_M + b_M/M. \quad (13)$$

Reiter’s derivation of T_p [41] used PPD but, subsequently, Reiter and Kinney [25] showed that PPD was not required for what they termed “partial synthesis” and that the estimator T_p could be used in this case. Their argument was based on inference for a Normal mean from completely synthesised data, as well as results from a simulation with incompletely synthesised data.

That the estimator T_p is also valid without PPD can be seen by noting that dropping the PPD step removes V from equations (11) and (12) and thus substituting (11) into (12) leaves equation (13) for T_p unchanged as an unbiased estimator of $var(Q|D_{syn})$. Reiter and Kinney argue that PPD is not needed here because synthetic populations are not generated for their examples, and only the observed units are retained in the synthetic data. They term this “partial synthesis”. We have shown here that this is not valid and the criterion for when PPD is needed is that different methods of inference are used for the real and synthetic data. This implies that T_p can also be used to estimate variances of completely synthesised data when this condition holds. When the sample sizes differ for observed and synthetic data equation (13) needs modification as

$$T_p = \bar{v}_M k/n + b_M/M. \quad (14)$$

Drechsler [8] has proposed this estimator for “fully synthetic data”, with sampling from the PPD, when both the observed and synthetic data are generated by SRS. For completely synthesised data the estimator T_s differs from T_p by replacing b_M/M with $(1 + k/n)\bar{v}_M$ for synthesis with PPD and by $(k/n)\bar{v}_M$ without. Since \bar{v}_M will be based on many more degrees of freedom than b_M/M we would expect T_s to be a more precise estimate than T_p . This was confirmed by the results of simulations, one of which is reported in A.1. Another advantage of T_s is that it does not require multiple syntheses.

2.4 Synthesising data from complex samples

In [39] the original development of methods for synthetic data assumed that synthetic data would be produced by SRS, whatever the design of the original survey data that were being synthesised. This approach would relieve the analyst with synthetic data of the need to use specialised methods for complex samples for their analyses. We will refer to such methods as complex-to-simple synthesis. In section 2.2 we have discussed the case when both the original and synthetic data are generated by simple random sampling, the simple-to-simple case. Agencies providing synthetic data from complex surveys may wish to provide data that can be analysed by the same methods as the observed data: the complex-to-complex case.

Reiter [40] used the complex-to-complex approach in the first paper to evaluate the variance estimator (T_M here, but T_s in [40]). The theory described in sections 2.1 to 2.3 will apply equally well to samples generated by complex sampling methods, provided the asymptotic assumptions outlined in 2.1 apply. If the same complex sampling methods are used to analyse the observed and synthetic data then the new simple variance estimators, T_s and $T_{s(PPD)}$, can be used and multiple syntheses will not be required. There are three aspects of the sample design that must be accommodated when analysing a complex sample 1) stratification, 2) weighting and 3) clustering, often in combination. We describe briefly below how these three factors affect the practical task of creating synthetic data.

2.4.1 Stratification

This is the easiest case. The sample is divided into strata and synthesis carried out within each stratum by the methods for SRS, if no weighting or clustering is involved, or by other appropriate methods otherwise. Reiter [40] presents a simulation to evaluate T_M and we extend this in A.2 to evaluate the properties of T_s and $T_{s(PPD)}$. Stratifying data before synthesis may also be of benefit, even when it is not required by the design, to preserve important interactions and to avoid computational problems with very large observed data sets.

2.4.2 Weighting

Weighting in sample surveys can come about for a number of reasons. In stratified samples different sampling fractions by stratum require inverse-probability weights. Other design features, such as the selection of only one individual in a household, may also mean that weighting is required to make inferences for the population of individuals. Weighting may also be used to adjust for non-response or post-stratification when the sample is matched to known population totals.

The weights will usually be assumed to be part of X , the data that will remain unchanged, although in some cases the weights themselves may be synthesised [34]. In either case, inference will always be conditional on the weights. A subgroup with weight w_i of size n_i will represent a proportion $w_i/\Sigma(n_i w_i)$ of the population and appropriate estimators for design weights are required for inference to the population. To create synthetic data for weighted samples synthetic samples are generated conditional on the values of the weights. This approach has been used by [52] in synthesising data from the American Community Survey adjusting for stratification, weighting and clustering. In general, this method will require that the relationship between the weights and other variables is correctly specified: an aspect of the ASD assumption. Where only stratification is involved then the weights will be constant within strata, so we can be certain that this aspect of the ASD assumption is met. If the ASD assumption is met for the data and their relationship to the weights, then analysis by methods that adjust for the design weights for both the observed and synthetic data justifies the use of the new estimators here.

We extend the simulation for stratified sampling in A.2 to the case of stratified weighted sampling and demonstrate that T_s and $T_{s(PPD)}$ give unbiased variance estimates in this case.

2.4.3 Clustering

Clustering can occur in sample surveys for one of two reasons. It may be introduced as part of the survey design, particularly for door-to-door surveys, as a means of reducing the cost of field work. Alternatively, clustering may be an aspect of the population being surveyed, such as the clustering of characteristics within small areas. Surveys with a cluster design are most easily analysed by the “ultimate-cluster” method, for example as implemented in the *vardpoor* package for **R** [5]. The ultimate-cluster method, which uses cluster-level summary statistics, would not provide a fit that could be used for synthetic data generation. However, the corresponding model-based analyses via hierarchical models can provide such a fit. This method has been used to generate synthetic data [40] and incorporated into a conditional approach by [52]. These applications have used PPD and the variance estimator T_M , but T_s and $T_{s(PPD)}$ could equally well be used if data are to be analysed as a complex sample. For the case where interest is in small geographic areas Drechsler and Hu [11] show that using geocoding data can preserve the geographic structure of the data without the need for defining clusters.

2.5 Summary and recommendations

The factors which determine which estimators should be used for the variances of estimates from synthetic data are 1) whether the data are completely or incompletely synthesised and 2) whether the methods of inference that would be used for the observed data match those to be used for the synthetic data. In sections 2.1 to 2.3 we dropped the conditioning on any unchanged variables X . If all of Y is replaced this is considered as complete synthesis for inferences that are conditional on X . Also, most importantly, all of the results for inference from synthetic data will only hold if the ASD assumption is true. Except for simulation studies, the ASD assumption is highly unlikely to be exactly true so all inference from synthetic data must be considered approximate. Based on the results above we make the following recommendations for the methods and estimators to use for synthetic data. We consider only the utility of synthetic data, independent of the impact of methods on the potential for disclosure, which is discussed in section 4.

The decision as to whether the data must be synthesised from the PPD of Y_{obs} depends on whether the methods of inference from the synthetic data are the same as those that would be used to analyse the observed data. Only if this is not the case, is PPD required. Using PPD when it is not required will result in increasing the variance of estimates from the synthetic data. Multiple synthetic data sets ($M > 1$) are necessary to obtain variance estimates when different methods are used for inference from the observed and synthetic data and also for incompletely synthesised data. Increasing M will reduce the variance of estimates from synthetic data, but only to a lower limit of U , the expected variance of Q estimated from the observed data.

For completely synthesised data when the methods of inference for the real and synthetic data match, the estimator T_s is the best choice. When this is not the case then T_M is the only choice, despite its relatively poor performance. If the original data were generated from a complex sample, and the synthetic data by SRS then information on the design factors, often published in technical reports, may allow U to be approximated from V . If this is the case PPD and multiple syntheses would not be required and equation (10) modified to give

$$T_{sDE} = \bar{v}_M(DEk/n + 1/M) \quad (15)$$

where DE is the approximate design effect. Although this is obviously approximate, it may be more useful in practice than using T_M . For incompletely synthesised data the estimator T_p should be used.

3 Practical aspects of data synthesis

The aim of the methods, as described in the literature we reviewed above, has been to provide synthetic data that can be used for inference to population parameters. In section 2 we have also adopted this approach, but it is not what we envisage for users of the LSs, at least at present. This is also true for users

of synthetic data products produced by the U.S. Bureau of the Census, where the publication of results from synthetic data alone is strongly discouraged. Staff with access to the original confidential data create synthetic data which are made available to analysts. The analysts carry out data preparation, exploratory analyses and preliminary modelling on the synthetic data. The code developed by the analysts is used to run final fits of candidate models and validation analyses on the original un-synthesised data. Concern about the validity of the ASD assumption underlies these policies. If this validation proves inadequate, the results can be used to improve the methods to be used for future clients. The term “gold-standard” describes the final steps where code for fitting and validating models is run on the actual data.

When this approach is being used the analysts will not be interested in any of the variance estimates discussed in 2. Instead, they will want an estimate of the results they might get from analysing the real data. This implies that synthetic data must be produced that matches the design of the real data, i.e. simple-to-simple or complex-to-complex. The analysts will want estimates, \bar{q}^M as before, and an estimate of the variance they might get from the real data which is provided by \bar{v}^M . One synthetic data set without sampling from the PPD will be all that is required in this case. The other estimators discussed here may come into their own if the verification process shows that synthetic data can be relied upon to give results that agree with those from the real data.

Real survey, census or administrative data may bear very little resemblance to the models used to derive the theory of synthetic data. Continuous data may have distributions that are nothing like a Normal distribution, even after a suitable transformation. Categorical data may have many complex interactions that it would be unreasonable to expect the synthesiser to investigate. Furthermore, real data may be subject to constraints that must be respected in the synthetic data. For example, if an analyst were presented with synthetic data on children in families where a natural parent was less than 12 years older than a child, the utility of the data would be questionable. Fortunately these questions have been addressed for synthetic data and the literature contains many options, some of which have been implemented in the *synthpop* package. Synthesis methods have been developed that preserve the marginal distributions of continuous variables [62] and these can be adapted to include an element of smoothing to prevent the identification of unique values. The use of a sequence of conditional distributions makes it easy to incorporate constraints on data values for synthetic samples. Variables that define the constraints must be synthesised first and the constrained variable is then synthesised with the constraint satisfied.

A number of methods from machine learning have been used successfully to generate synthetic data [6, 14, 44]. Classification and regression tree models (CART) performed well in the evaluation carried out by Drechsler and Reiter [14]. These have been implemented in *synthpop* with a choice of algorithms from the **R** packages *party* [23] or *rpart* [55]. We have found that these methods are able to reproduce the main features of data sets without the need for exploratory analysis. One such example is presented in section 5. This has led us to set CART as the default method in *synthpop* if no detailed models are specified by the analyst. Such models may seem distant from the theoretical models used when describing the theory of synthetic data. But the adaptive nature of a series of conditional CART models may give a better approximation to the structure of the observed data than is possible within the constraints of parametric models. They may come under George Box’s category [3] of being “wrong but useful”.

When data to be synthesised have missing values they could be replaced by imputations and the multiply imputed data sets can then be synthesised. This has been implemented [2, 9] using formulae and variance estimates for combining multiple imputations with multiple syntheses [42]. We have not adopted this approach in *synthpop* because we expect that the choices about handling missing data for a particular project should be the responsibility of the analyst. Using a missing-at-random approach we synthesise the missingness indicator first, and then synthesise the remaining cases from a fit to the non-missing cases in the observed data. Both the synthesised values and the missingness indicator can then be used together in the synthesis of later variables. This guarantees that any relationships with the missingness indicator are maintained in the synthetic data. For variables earlier in the sequence this is assured by including them as predictors of missingness and, for those later in the sequence, by having the missingness indicator as one of their predictors. An analyst can use synthesised data with missing values to decide how to handle them and their methods can readily be run on the real data. In some

cases further synthetic data, with missing values ignored or imputed, could be provided to the analyst.

Data from the LSs often includes time to event data. These may be defined as a series of dates or as a follow-up time and an indicator of the event at the end of follow up. For example, the LSs are linked to death registrations and emigration records. To synthesise such data the event indicator is synthesised first and the follow-up times are synthesised separately for each type of event. Possible models for follow-up times are parametric survival models (Weibull or log-Normal) or a CART method applied to Kaplan-Meier survival estimates [23]. Cox proportional hazard models would be more difficult to fit because they would require the vector of all observed event times and the corresponding baseline hazard to be used in simulating the synthetic data. Poisson models can be used for person-years analyses.

4 Utility and disclosure risk for synthetic data

Analysis of synthetic data will never yield exactly the same results as would be found from the observed data. We hope synthetic data will be useful and the term “utility” is used for the extent to which results from synthetic data agree with those from the original data. Utility measures may be specific to the inferences being made from the data, as in Figures 3 and 4 in Section 5, or they may be general utility measures that assess the overall difference between the observed and synthesised distributions. In both cases the utility of synthetic data will depend crucially on the ASD assumption. Proposals for general utility measures have been suggested by Woo et al. [61] and are adapted for the specific case of synthetic data by Snoke et al. [54].

Although completely synthesised data contain no records that are identified with real cases, this does not mean that their real or perceived disclosure risk is zero. Fienberg et al. [20] have proposed a Bayesian method of assessing disclosure risk by modelling optimal intruder behaviour and this methodology has been employed to assess disclosure risk for synthetic data [12, 46]. Similar methodology has been used by Elliot [18] to evaluate output from the *synthpop* package. These methods are specific to the type of attack that an intruder is presumed to be attempting. The development of ϵ -differential privacy [15] provides a measure of disclosure risk which does not depend on the specific tactics of an intruder, although it assumes the intruder possesses a very large amount of ancillary knowledge about the real data. An extension to this method, ϵ - δ -differential privacy [16, 17], which allows a mechanism (e.g. a synthesis) to satisfy ϵ -differential privacy with probability $1 - \delta$, has been used to generate the synthetic data that sit behind the U.S. Bureau of the Census’s “OnTheMap” application², although the synthesising algorithm required adjustments to yield useful results [31]. It has been shown [1, 7] that the use of this method of creating ϵ - δ -differentially-private synthetic data, without adjustment, invalidates the usual methods of inference from synthetic data. More recent research with synthesised categorical data [24, 32, 33, 48] has proposed methods that make some of the same assumptions as differential privacy and can be used to identify individual records with a high disclosure potential.

How will our recommendations for the analysis of completely synthesised data affect any disclosure risk posed by the synthetic data? By not sampling from the PPD we are introducing less noise than would be the case with PPD. But the need to generate only a single synthetic data set should reduce the disclosure risk considerably. Indeed concern about a potential disclosure risk has led agencies to permit the release of only a single synthetic data set [22]. Our proposals to supply synthetic data to users had to be agreed by the SLS Research Board and we have not asked for permission to release more than one synthetic data set to a user. The report mentioned above [18] formed part of our submission to the SLS Research Board, but another aspect of disclosure control was of concern to them. If an intruder sees the synthetic data they may mistakenly believe it to be real and attempt an identification, with subsequent loss of reputation for the data collection agencies. We have agreed to carry out additional disclosure control measures to avoid any such problem. These are now incorporated into the *synthpop* package. All data sets have to be clearly labelled as “false data”. Any observations that are unique in the actual data and also appear as unique in the synthetic data have to be removed from the latter. Options for top-coding of outlying observations are also available as part of the disclosure control function in

²U.S. Census Bureau. 2015. OnTheMap Application. Longitudinal-Employer Household Dynamics Program. <http://onthemap.ces.census.gov/>

synthpop.

The computer science literature uses the idea of the disclosure-utility trade-off, e.g. [29, 53], implying that increasing utility will also increase disclosure risk. This need not be the case for synthetic data since the use of a more appropriate model may increase utility without necessarily increasing disclosure risk. Ideally, an analyst creating synthetic data should be able to obtain interpretable measures of utility and disclosure risk to evaluate the synthesis. In the *synthpop* package we have implemented this for utility, but a measure of disclosure risk is not currently available for data with the complexity of most survey and administrative data. We hope that further research in this area will help us to meet this challenge in future.

5 Example

5.1 Methods

The fitting of formal models is only a small part of any statistical analysis. The majority of an analyst’s time is taking up with checking and exploring the data and in carrying out preliminary tabulations. We aim to produce synthetic data that can be used for this type of analysis. To illustrate synthesis of LS data we have extracted data on age, sex, marital status, ethnic group and long-term illness from the SLS database for the 1991 and 2001 Censuses. The acronyms AGE9, SEX9, MSTAT9, ETH9, ILL9 and AGE0, SEX0, MSTAT0, ETH0, ILL0 are used to describe them. The synthesis was carried out for over 186 thousand SLS members who were present at both censuses. No preliminary data cleaning was carried out on the extract. Some variables had a small percentage of missing values. For the categorical variables the missing cases are simply handled as an additional category, but AGE0 had 0.15% of missing cases coded as -999 and this formed a missingness indicator for AGE0. Exploratory analyses of univariate distributions, cross tabulations and results from fitted models were compared for the real and the synthetic data. The *synthpop* package was used for all analyses and more details of the methods can be found in [35].

Initial synthesis was carried out without PPD for two choices of models. The first (CART) used CART for all the variables and the second (Parametric) used an appropriate parametric method for each variable: polychotomous or logistic regression for categorical data and distribution-preserving linear regression for AGE9 and AGE02. A lower triangular prediction matrix was used in both cases so that all previously synthesised variables were used in the predictions. The ordering of the variables used in the final syntheses and for the results presented here was ETH9, ETH0, AGE9, SEX9, MSTAT9, ILL9, AGE0, SEX0, MSTAT0, ILL0.

5.2 Results of exploratory analyses

A problem with the initial run of the synthetic data was that some SLS members under 16 had marital status “married” in the synthetic data, with the number of such cases being larger for the Parametric syntheses. This was readily fixed by imposing a logical rule during the syntheses. Marginal distributions of all the variables were comparable to those for the real data for syntheses produced from CART, but for parametric methods, synthetic data for two variables had marginal distributions that did not match the observed data. Initially the variables ETH9 and ETH0 were placed later in the synthesis order. This produced unsatisfactory results for small ethnic groups, when parametric methods were used. The most extreme differences were for Bangladeshis, the smallest category with only 28 SLS members in 2001, who were over-represented with an average of over 4,000 in the synthetic data sets. Exploratory analyses showed that these problems arose from attempting to fit a polychotomous regression to sparse data with many zero cells, when some parameters are fitted at their extremity. Moving the two ethnic group variables to the start of the Parametric synthesis overcame the problem. Another approach of defining a predictor matrix with smaller models used to predict the ethnic group variables also gave satisfactory results.

The other variable that was affected was AGE0 (see Figure 1). From Figure 2 we can see why




Figure 1: Comparison of real (black) and synthetic (grey) data for AGE0 in 5 year age groups




Figure 2: Plots of AGE9 against AGE0 for real data and one synthesis by each of CART and parametric methods (most points refer to multiple observations)

the parametric method failed to reproduce the distribution of AGE0³. Age is recorded in full years and, in most cases, AGE0 is exactly AGE9+10. Exceptions could be those with birthdays between the dates of the two censuses as well as various data errors or mismatches. Differences of exactly 10 or 20 years are common. The parametric method could not reproduce this pattern. The parametric syntheses were rerun with the method for just this variable changed to CART giving satisfactory results. Thus for these exploratory analyses the CART method gives more satisfactory results than parametric methods, and with no requirement to customise the analyses in any way.

5.3 Results of fitting models to data synthesised without PPD

We now investigate how inference from synthetic data differs from that from the real data using variables from the data set described in the last section. A logistic regression modelled absence of long-term illness in 1991 (ILL9) from AGE9, MSTAT9 and SEX9. We use the same two synthesis models (Parametric and CART) described above with $M = 10$ in each case, the rule for marital status for the under 16s in place and ILL9 as the last variable synthesised. Initially a model with the main effects of each variable was

³Points at ages over 90 are not shown to avoid any possible disclosure of extreme ages.




Figure 3: Coefficients of fit to ILL9=“No” from AGE9, SEX9 and MSTAT9 for real and synthetic data

fitted to the real and the synthetic data sets. Figure 3 compares the estimates from the real data with the averages from the 10 simple syntheses from parametric and CART models. Confidence intervals from the synthetic data sets are calculated from \bar{v}_M , the estimate of the variance that would be expected from the observed data. Since the final step in the synthesis was a logistic regression of ILL9 on these three variables, we know that the ASD assumption will be met for this model. As expected, the estimates from the parametric model are very close to the real estimates, and a formal test showed that there was no evidence of any bias. In this case the estimates from the CART syntheses are also fairly close and certainly would be satisfactory for exploratory analyses. We can see that freedom from long-term illness decreases sharply with age and is higher for females than males. Adjusting for age and sex, those married, remarried or widowed are more likely to be free from long-term illness than those who are single, whereas the opposite is true for the divorced.

A further model was fitted which includes a sex by marital status interaction, so that the ASD assumption will not hold for the synthesis from Parametric models. Results are shown in Figure 4. For the real data there is evidence of an interaction. The association of being married with lack of illness is stronger for men than for women. The CART syntheses do a reasonable job of reproducing this, whereas the parametric syntheses show no evidence of this interaction since they are generated from an interaction-free model. The analyst with access only to the synthetic data generated by this parametric model would find no evidence of this interaction in any diagnostic plots. The CART model may be less exact when the parametric model is correct, but its flexibility shows evidence of an interaction from the synthetic data.

These results all assume that the analyst is interested in estimating the results which might be obtained from the real data, rather than in making inferences to population parameters directly from the synthetic data. We believe this is an appropriate use of synthetic data because we can never be completely sure that the ASD condition is met. Further experience with synthetic data may show that this is too cautious a position, but it is a safe one for now.




Figure 4: Coefficients of fit to ILL9=“No” from AGE9, SEX9 and MSTAT9*SEX9 interaction for real and synthetic data

5.4 Inference to population parameters from data generated with and without PPD

Finally, we present inferences for the population parameters for this example to illustrate the properties of the variance estimates described above. We fitted the same model as in the previous section, with just the main effects, with parametric and CART methods and compared variance estimates with those from the observed data. CART synthesis with PPD starts by taking a bootstrap sample of the real data and fitting the CART models to this sample. For synthesis without PPD we have a single estimate T_s whereas, for synthesis with PPD, we have three variance estimates $T_{s(PPD)}$, T_p and T_M . Proper syntheses with parametric and CART models were carried out with $M = 10$ in each case.

The synthetic estimates of the coefficients all showed the same patterns for CART and parametric methods as in Figure 3. The expected values of the variances of \bar{q}_M as an estimate of Q should be $V(1 + 1/M)$ without PPD and $V(1 + 2/M)$ with PPD, since $k = n$. Table 2 gives the ratios of the standard errors divided by \sqrt{V} , estimated from the standard error from the observed data. If the ASD assumption holds the expected value of these ratios of standard errors would be approximately 1.049 without PPD and 1.095 with PPD. For the CART syntheses, we can see that standard errors calculated from T_s and $T_{s(PPD)}$ are close to their expectations. Those calculated from T_p are somewhat larger for the CART methods. The results are more irregular for the parametric methods which gave more biased estimates of the standard errors. In both cases, as expected, the estimator from T_M is quite unsatisfactory, giving a negative value in one case.

These results support our recommendation of the use of the variance estimates T_s for completely synthesised data.

6 Summary and future directions

It is now over twenty years since the first proposals were made to use synthetic data for disclosure control. A rich literature has since been developed, mainly by Reiter and his colleagues on the theory and practice of synthetic data generation. We were able to draw on this in developing the *synthpop* package. While we make some suggestions for different interpretations of the methods and introduce

Table 1: Standard errors for estimating the population coefficients of a logistic regression of ILLP9 predicted from AGE9, SEX9 and MSTATT9 synthetic data produced by parametric and CART methods, with and without PPD. Standard errors are calculated from different variance estimates and expressed as proportions of the standard errors from the observed data

<i>Ratios to standard error from the real data</i>								
	<i>Parametric</i>				<i>CART</i>			
	<i>without PPD</i>	<i>with PPD</i>		$\sqrt{T_M}$	<i>without PPD</i>	<i>with PPD</i>		$\sqrt{T_M}$
	$\sqrt{T_s}$	$\sqrt{T_{s(PPD)}}$	$\sqrt{T_p}$		$\sqrt{T_s}$	$\sqrt{T_{s(PPD)}}$	$\sqrt{T_p}$	
Intercept	1.046	1.091	1.078	0.940	1.045	1.098	1.107	1.201
AGE9	1.165	1.217	1.196	0.952	1.048	1.095	1.133	1.458
SEX9								
Female	1.043	1.089	1.174	1.820	1.049	1.095	1.131	1.447
MSTAT9								
Married	1.172	1.225	1.248	1.465	1.048	1.097	1.232	2.159
Remarried	1.086	1.135	1.103	0.704	1.052	1.095	1.119	1.337
Divorced	1.091	1.143	1.081	NA	1.064	1.089	1.212	2.076
Widowed	1.145	1.199	1.208	1.294	1.048	1.099	1.216	2.044

NA Negative variance estimate

some new variance estimators, our work depends heavily on their insights.

There is much more we still need to learn about the best way to carry out syntheses. Recommendations are needed on choosing the ordering of variables during synthesis, for deciding whether reduced models excluding some variables should be used and for fine-tuning the parameters of CART models. The *synthpop* package is intended to facilitate this and new methods can be added by the user that are not at present part of the package. The synthesiser, with access to the real data, ought to carry out checks on the validity of the data before it is released to the analyst. At a minimum, a visual check on all the marginal distributions should be carried out. The *synthpop* package includes a function to do this and Figure 3 is an example of part of its output. Now that registered and accredited users of the SLSs are being supplied with synthetic data, their feedback on its performance will help us to develop best practice. We encourage users to include code to verify any assumptions (such as the absence of interactions) when a gold-standard analysis is run.

The structure of administrative data also presents challenges. In particular, we hope to develop methods for repeat events, such as hospital admissions and for the synthesis of data for all members of a family. Options to include utility measures are already being provided in *synthpop*, but the ability to include measures of disclosure risk remains a more distant goal.

A Appendix

A.1 Simulation to evaluate methods for synthetic data with simple random sampling

This simulation is similar to one used to evaluate methods for synthetic data and the variance estimator T_M in [39]. The population was created by drawing a sample of size $N = 50,000$ from $N(0, \Sigma)$, where Σ is a 5 by 5 matrix with diagonal elements 1 and off diagonals 0.5. The columns of the population are denoted by $(y_1, y_2, y_3, y_4, y_5)$. Each of 10,000 simulations selected a random sample of size $n = 500$ from this population as the observed data. The parameters of the five-variate Normal distribution were estimated for each simulation and two sets of synthetic data were generated, each with $M = 5$ replicates of size $k = 1,000$, one set with PPD and one without. For the observed data the fpc is close to unity here, at 0.99. Simulations where this was not the case gave estimated variances that were biased, illustrating

the need for this assumption to be met, as discussed in 2.1.

A linear model for y_1 predicted from the other 4 variables was used to evaluate methods of inference. The simulations without PPD confirmed that T_s and T_p are unbiased estimates of the variance of the coefficients. The same was true for T_M , $T_{s(PPD)}$ and T_p from the simulations with PPD. But T_M was only unbiased because of the inclusion of approximately 11% of results with negative variances, which would never be accepted in practice. Two approaches were used to adjust the negative values of T_M ; basing the intervals on only the cases where $T_M > 0$ or using an estimator T_M^{adj} , proposed by [40]. In both these cases the variance estimates were biased upwards by 13% and 10% respectively.

Table 2: Ratio of the variance from the simulations for existing estimators compared to the variance of the new estimators T_s and $T_{s(PPD)}$

	without PPD	with PPD			
	ratio to variance of T_s	ratio to variance of $T_{s(PPD)}$			
	T_p	T_p	T_M	T_M^{adj}	$T_M > 0$
Intercept	1.94	5.68	173.44	145.47	161.97
y2	1.49	3.54	94.77	81.13	89.19
y3	1.45	3.52	94.27	80.25	88.61
y4	1.45	3.52	94.14	86.3	85.8
y5	1.48	3.49	93.65	80.15	88.01

The variance, estimated from the simulations, of T_M (and its adjusted variants) and T_p , relative to the new simple estimates are shown in Table 2. The very poor precision of all the T_M estimators is evident. The new simple estimators T_s and $T_{s(PPD)}$ have better precision than T_p . Table 3 shows the estimated coverage calculated for all these variance estimates, except for T_M , where the negative values prevent the estimate of coverage. In practice, an analyst faced with a negative variance estimate, would either abandon the whole enterprise and perhaps ask for a new set of synthetic data, or use the adjusted estimator T_M^{adj} . We can see that all the results are very satisfactory except for the adjusted T_M estimators, which have reduced coverage.

Table 3: Coverage of 95 % confidence intervals for Q calculated from different variance estimators

	without PPD			with PPD		
	T_s	T_p	$T_{s(PPD)}$	T_p	T_M^{adj}	$T_M > 0$
intercept	94.7	94.7	95.0	94.7	86.1	85.4
y2	95.4	95.3	95.0	94.8	96.8	86.3
y3	95.1	95.2	95.2	94.9	86.5	85.7
y4	94.9	94.9	94.8	94.4	86.3	85.8
y5	95.3	95.2	95.0	94.7	86.6	86.1

A.2 Simulation to evaluate synthesis for a stratified sample

A.2.1 Methods

This simulation was based on that presented in section 3.2 of [40] where the simulated observed data and the synthetic data are drawn as stratified samples. A population of size $N = 10,000$ was created consisting of 10 strata $h = 1, 2, \dots, 10$ each of size $N_h = 1,000$ and where the variable Y of interest is distributed as $N(10h, h^2)$, in the h^{th} stratum. Each simulation selected a stratified random sample of total size $n = 200$, with n_h from each stratum, from this population to represent the observed data.

Synthetic samples of size 200 were then generated. In [40] this was done by first taking a random sample from the population without constraining it to balance the population totals. In the simulation reported here exactly $n_h = 20$ values were synthesised for each stratum so that inference from the synthetic data will be the same as from the observed data. The information on stratum membership, which is fixed for each synthetic sample is the X_{obs} for this example, and all inference is conditional on X_{obs} .

The parameter (Q) to be estimated is the population mean estimated by the stratified estimate $\sum_{h=1}^{10} (N_h/N) \bar{y}_h$ where \bar{y}_h is the mean for of the observations in the h^{th} stratum with variance estimated from (1).

$$\sum_{h=1}^{10} \left(1 - \frac{n_h}{N_h}\right) \left(\frac{N_h}{N}\right)^2 \frac{\sum_{j=1}^{n_h} (y_{hj} - \bar{y}_h)^2}{(n_h - 1) n_h} \quad (16)$$

In [40] the value of n_h was 20 for all h , so no weighting is involved, and $M=100$ synthetic data sets were produced for each simulation. This sample design has an extreme design effect. If the usual SRS formula for the variance of the mean were to be used the result would have a variance more that 22 times greater than if the correct one were used. We have extended the simulation to a more realistic number of syntheses ($M = 10$) and to different sampling fractions in each stratum. For each of the following three cases we ran 1,000 simulations of synthesis with and without PPD.

1. Unweighted simulation with $M = 100$ and all $n_h = 20$, as in [40].
2. As the first simulation but with $M = 10$.
3. Selecting samples of size 11, 13, 15, 17, 19, 21, 23, 25, 27, 29 from strata 1 to 10. This gives weights that range from 91 to 34. The larger sampling fraction in the more variable strata should give a reduced variance, although to a lesser extent than the optimal sampling scheme which would give much more extreme weights.

A.2.2 Results

All the simulations gave unbiased estimates of Q . For synthesis with PPD simulations 2 and 3 had 6.5% and 6.2% of results with negative estimates for T_M , while there were no negative estimates for simulation 1 with $M = 100$. For simulations 2 and 3, only the adjusted estimates from T_M could be used for confidence intervals.

Table 4 compares the variances from the simulations with the mean of the different estimators as well as the coverage of nominal 95% intervals for based on each. The variance estimate T_M from proper synthesis is slightly biased upwards, as was found in the original publication. This may relate to the use of the fpc which would be valid for the observed data, but not appropriate for the synthetic data. All the other estimators appear unbiased, at least to the accuracy provided by the simulation. The adjustments to T_M required in simulations 2 and 3 resulted in increased variances and lower confidence interval coverage. The coverage of intervals based on all other variance estimators appears satisfactory.

References

- [1] ABOWD, J. M., AND VILHUBER, L. How protective are synthetic data? In *Privacy in Statistical Databases, Lecture Notes in Computer Science, PSD 2008* (2008), J. Domingo-Ferrer and Y. Saygin, Eds., Springer Berlin Heidelberg, pp. 239–246.
- [2] BENEDETTO, G., STINSON, M. H., AND ABOWD, J. M. The creation and use of the SIPP Synthetic Beta. https://www.census.gov/content/dam/Census/programs-surveys/sipp/methodology/SSBdescribe_nontotals.pdf 2013.
- [3] BOX, G. E. P. Science and statistics. *Journal of the American Statistical Association* 71 (1976), 791–799.

Table 4: Variance of the estimates of the mean from stratified sampling from the simulations compared to the average of the different variance estimators and the coverage of the corresponding of 95% confidence intervals

	Simulation 1		Simulation 2		Simulation 3	
	Variance	coverage	Variance	coverage	Variance	coverage
Without PPD						
From simulation	0.200		0.220		0.166	
T_s	0.189	94.1	0.206	94.9	0.173	94.5
T_p	0.189	94.1	0.206	94.5	0.174	94.5
With PPD						
From simulation	0.197		0.242		0.191	
$T_{s(PPD)}$	0.191	94.6	0.225	93.9	0.189	94.5
T_p	0.191	94.7	0.229	93.8	0.193	94.8
T_M	0.244	95.7	0.279		0.237	
T_M^{alt}	0.244	95.7	0.294	88.9	0.249	90.1
$T_M > 0$	0.244	95.7	0.302	88.8	0.256	90.1

- [4] BOYLE, P., FEIJTEN, P., FENG, Z., HATTERSLEY, L., HUANG, Z., NOLAN, J., AND RAAB, G. M. Cohort Profile: the Scottish Longitudinal Study (SLS). *International Journal of Epidemiology* 38 (2009), 385–392.
- [5] BREIDAKS, J., LIBERTS, M., AND IVANOVA, S. Variance estimation for sample surveys by the ultimate cluster method. Manual, <https://cran.r-project.org/web/packages/wardpoor/wardpoor.pdf>, 2015. Accessed: 2016-01-05.
- [6] CAIOLA, G., AND REITER, J. P. Random forests for generating partially synthetic, categorical data. *Transactions on Data Privacy* 3 (2010), 27–42.
- [7] CHAREST, A. How can we analyze differentially-private synthetic data sets. *Journal of Privacy and Confidentiality* 2 (2010), 21–33.
- [8] DRECHSLER, J. Improved variance estimation for fully synthetic datasets. Joint UN-ECE/Eurostat work session on statistical data confidentiality 2011. Working paper, 18. http://www.unece.org/fileadmin/DAM/stats/documents/ece/ces/ge.46/2011/18_Drechsler.pdf, 2011. Accessed: 2014-07-14.
- [9] DRECHSLER, J. *Synthetic Data Sets for Statistical Disclosure Control*. Springer, New York, 2011.
- [10] DRECHSLER, J. New data dissemination approaches in old Europe - synthetic datasets for a German establishment survey. *Journal of Applied Statistics* 39 (2012), 243–265.
- [11] DRECHSLER, J., AND HU, S. Generating synthetic geocoding information for public release. Joint UNECE/Eurostat work session on statistical data confidentiality 2015. Available from <http://www1.unece.org/stat/platform/display/SDCWS15>, 2015. Accessed: 2015-12-14.
- [12] DRECHSLER, J., AND REITER, J. P. Disclosure risk and data utility for partially synthetic data: An empirical study using the German IAB establishment survey. *Journal of Official Statistics* 25 (2009), 589–603.
- [13] DRECHSLER, J., AND REITER, J. P. Sampling with synthesis: a new approach for releasing public use census microdata. *Journal of the American Statistical Association* 105 (2010), 1347–1357.

- [14] DRECHSLER, J., AND REITER, J. P. An empirical evaluation of easily implemented, nonparametric methods for generating synthetic datasets, computational statistics and data analysis. *Computational Statistics and Data Analysis* 55 (2011), 3232 – 3243.
- [15] DWORK, C. Differential privacy. In *Automata, Languages and Programming: 33rd International Colloquium, ICALP 2006, Proceedings, Part II* (2006), M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener, Eds., Springer Berlin Heidelberg, pp. 1–12.
- [16] DWORK, C., KENTHAPADI, K., MCSHERRY, F., MIRONOV, I., AND NAOR, M. Our data, ourselves: Privacy via distributed noise generation. In *Advances in Cryptology - EUROCRYPT 2006: 24th Annual International Conference on the Theory and Applications of Cryptographic Techniques* (2006), S. Vaudenay, Ed., Springer Berlin Heidelberg, pp. 486–503.
- [17] DWORK, C., AND SMITH, A. Differential privacy for statistics: What we know and what we want to learn. *Journal of Privacy and Confidentiality* 2 (2011), 135–154.
- [18] ELLIOT, M. Final report on the disclosure risk associated with the synthetic data produced by the SYLLS team. <http://www.cmist.manchester.ac.uk/research/publications/reports/>, 2014.
- [19] ELLIOT, M., AND PURDAM, K. A case study of the impact of statistical disclosure control on data quality in the individual uk samples of anonymized records. *Environment and Planning A* 39 (2007), 1101–1118.
- [20] FIENBERG, S. E., MAKOV, U. E., AND SANIL, A. P. A Bayesian approach to data disclosure: optimal intruder behavior for continuous data. *Journal of Official Statistics* 13 (1997), 75–89.
- [21] HATTERSLEY, L., AND CRESSER, R. *The Longitudinal Study, 1971-1991: History, organisation and quality of data. LS Series no.7*. London: The Stationery Office, 1995.
- [22] HAWALA, S. Partially synthetic data via expert knowledge modeling, and matching. In *SSC Annual Meeting, May 2008, Proceedings of the Survey Methods Section* (2008). Available at http://www.ssc.ca/survey/documents/SSC2008_S_Hawala.pdf. Accessed 2016-01-12.
- [23] HOTHORN, T., HORNIK, K., AND ZEILEIS, A. party: A laboratory for recursive partytioning. R package version 1.0-13. <https://cran.r-project.org/web/packages/party/vignettes/party.pdf>, 2011.
- [24] HU, J., REITER, J. P., AND WANG, Q. Disclosure risk evaluation for fully synthetic data. In *Privacy in Statistical Databases, Lecture Notes in Computer Science, PSD 2014* (2014), J. Domingo-Ferrer, Ed., vol. 8744, Springer International Publishing, pp. 185–199.
- [25] KINNEY, S. K., AND REITER, J. P. Tests of multivariate hypotheses when using multiple imputation for missing data and disclosure limitation. *Journal of Official Statistics* 26 (2010), 301–315.
- [26] KLEIN, M., AND SINHA, B. Inference for singly imputed synthetic data based on posterior predictive sampling under multivariate normal and multiple linear regression models. *Sankhya B* 77, 2 (2015), 293–311.
- [27] KLEIN, M., AND SINHA, B. Likelihood based finite sample inference for singly imputed synthetic data under the multivariate normal and multiple linear regression models. *Submitted* (2015).
- [28] KLEIN, M., AND SINHA, B. Likelihood-based inference for singly and multiply imputed synthetic data under a normal model. *Statistics and Probability Letters* 106 (2015).
- [29] LOUKIDES, G., GKOUALAS-DIVANIS, A., AND SHAO, J. On balancing disclosure risk and data utility in transaction data sharing using R-U confidentiality map. Joint UNECE/Eurostat work session on statistical data confidentiality https://www.unece.org/fileadmin/DAM/stats/documents/ece/ces/ge.46/2011/19_Goulalas-Divanis_et_al. 2011.

- [30] LUMLEY, T. *Complex Surveys: A Guide to Analysis Using R*. Wiley, Hoboken, New Jersey, 2010.
- [31] MACHANAVAJJHALA, A., KIFER, D., ABOWD, J., GEHRKE, J., AND VILHUBER, L. Privacy: Theory meets practice on the map. In *International Conference on Data Engineering, ICDE 2008* (2008), pp. 277–286.
- [32] MCCLURE, D., AND REITER, J. P. Differential privacy and statistical risk measures: an investigation with binary synthetic data. *Transactions on Data Privacy* 5 (2012), 535–52.
- [33] MCCLURE, D., AND REITER, J. P. Assessing disclosure risks for synthetic data with arbitrary intruder knowledge. *Statistical Journal of the International Association for Official Statistics* (Submitted). Available from <http://www.ncrn.info/bibliography/assessing-disclosure-risks-synthetic-data-arbitrary-intruder-kn>
- [34] MITRA, R., AND REITER, J. P. Adjusting survey weights when altering identifying design variables via synthetic data. In *Privacy in Statistical Databases, Lecture Notes in Computer Science, PSD 2006* (2006), J. Domingo-Ferrer and L. Franconi, Eds., Springer Berlin Heidelberg, pp. 177–188.
- [35] NOWOK, B., RAAB, G. M., AND DIBBEN, C. synthpop : Bespoke creation of synthetic data in R. Package vignette <http://cran.r-project.org/web/packages/synthpop/vignettes/synthpop.pdf>, 2015. Accessed: 2015-02-26.
- [36] OHM, P. Broken promises of privacy: responding to the surprising failure of anonymization. *UCLA Law Review* 57 (2010), 1701–1775.
- [37] O’REILLY, D., ROSATO, M., CATNEY, G., JOHNSTON, F., AND BROLLY, M. Cohort description: The Northern Ireland Longitudinal Study (NILS). *International Journal of Epidemiology* 41 (2011), 634–641.
- [38] RAGHUNATHAN, T., LEPKOWSKI, J., VAN HOEWYK, J., AND SOLENBERGER, P. A multivariate technique for multiply imputing missing values using a sequence of regression models. *Survey Methodology* 27 (2001).
- [39] RAGHUNATHAN, T. E., REITER, J. P., AND RUBIN, D. B. Multiple imputation for statistical disclosure limitation. *Journal of Official Statistics*. 19 (2003), 1–17.
- [40] REITER, J. P. Satisfying disclosure restrictions with synthetic data sets. *Journal of Official Statistics* 18 (2002), 531–544.
- [41] REITER, J. P. Inference for partially synthetic, public use microdata sets. *Survey Methodology* 29 (2003), 181–188.
- [42] REITER, J. P. Simultaneous use of multiple imputation for missing data and disclosure limitation. *Survey Methodology* 30 (2004), 235–242.
- [43] REITER, J. P. Releasing multiply imputed, synthetic public use microdata: an illustration and empirical study. *Journal of the Royal Statistical Society: Series A* 168 (2005), 185–205.
- [44] REITER, J. P. Using CART to generate partially synthetic, public use microdata. *Journal of Official Statistics* 21 (2005), 441–462.
- [45] REITER, J. P., AND KINNEY, S. K. Inferentially valid, partially synthetic data: generating from posterior predictive distributions not necessary. *Journal of Official Statistics* 28 (2012), 583–590.
- [46] REITER, J. P., AND MITRA, R. Estimating risks of identification disclosure in partially synthetic data. *Journal of Privacy and Confidentiality* 1 (2009), 99–110.
- [47] REITER, J. P., AND RAGHUNATHAN, E. The multiple adaptations of multiple imputation. *Journal of the American Statistical Association* 102 (2007), 1462–1471.

- [48] REITER, J. P., WANG, Q., AND ZHANG, B. R. Bayesian estimation of disclosure risks in multiply imputed, synthetic data. *Journal of Privacy and Confidentiality* 6 (2014), 17–33.
- [49] ROYSTON, P. Multiple imputation of missing values. *Stata Journal* 4, 3 (2004), 227–241(15).
- [50] RUBIN, D. B. *Multiple Imputation for Nonresponse in Surveys*. Wiley, New York, 1987.
- [51] RUBIN, D. B. Discussion: Statistical disclosure limitation. *Journal of Official Statistics* 9 (1993), 462–468.
- [52] SAKSHAUG, J. W., AND RAGHUNATHAN, T. E. Generating synthetic microdata to estimate small area statistics in the american community survey. *Statistics in Transition* 15 (2014), 341–368.
- [53] SANKAR, L., RAJAGOPALAN, S. R., AND POOR, H. V. An information-theoretic approach to privacy. Proceedings of the 48th Annual Allerton Conference Communication Control and Computing http://lalithasankar.faculty.asu.edu/UP_DB_Allerton_ver1_2.pdf, 2011.
- [54] SNOKE, J., NOWOK, B., RAAB, G. M., DIBBEN, C., AND SLAVKOVIC, A. General and specific utility measures for synthetic data. In preparation, 2016.
- [55] THERNEAU, T. M., AND ATKINSON, E. J. An introduction to recursive partitioning using the RPART routines. <https://cran.r-project.org/web/packages/rpart/vignettes/longintro.pdf>, 2015.
- [56] US CENSUS BUREAU. Synthetic Longitudinal Business Database. Washington DC; Cornell University, Synthetic Data Server, Ithaca, NY, 2011. <http://www.census.gov/ces/dataproducts/synlbd/methodology.html>, 2011. Accessed: 2015-02-25.
- [57] US CENSUS BUREAU. SIPP Synthetic Beta: Version 5.1. Washington DC; Cornell University, Synthetic Data Server, Ithaca, NY, 2013. <http://www.census.gov/programs-surveys/sipp/methodology/sipp-synthetic-beta-data-product.html>, 2013. Accessed: 2015-02-25.
- [58] VAN BUUREN, S. Multiple imputation of discrete and continuous data by fully conditional specification. *Statistical Methods in Medical Research* 16 (2007), 219–242.
- [59] VAN BUUREN, S., AND GROOTHUIS-ODSHOORN, K. mice: Multivariate imputation by chained equations in R. *Journal of Statistical Software* 45 (2011), 1–67.
- [60] WINKLER, W. E. Examples of easy-to-implement, widely used methods of masking for which analytical properties are not justified. Technical Report Series 2007 21, Statistical Research Division, U.S. Census Bureau, Washington, DC, 20233, 2007.
- [61] WOO, M.-J., REITER, J. P., OGANIAN, A., AND KARR, A. F. Global measures of data utility for microdata masked for disclosure limitation. *Journal of Privacy and Confidentiality* 1 (2009).
- [62] WOODCOCK, S. D., AND BENEDETTO, G. Distribution-preserving statistical disclosure limitation. *Computational Statistics and Data Analysis* 53 (2009), 4228–4242.